

REMARKS

Claims 1-13 and 19-23 stand rejected. Claim 1 has been amended to set forth the recited subject matter more clearly. New claims 30-34 have been added. Reconsideration of the application in view of the remarks set forth below is respectfully requested.

Election/Restrictions

The Examiner restricted claims 1-13 and 19-23 to Group I and restricted claims 14-18 and 24-29 to Group II. Specifically, the Examiner restricted the claims of Group I as being drawn to handling deferred replies at a bus bridge (class 710, subclass 306) and restricted the claims of Group II as being drawn to request blocking (class 710, subclass 200). During a telephone conversation between prosecuting attorney of record, Robert Manware, and Examiner Trisha Vu, on March 19, 2004, a provisional election was made without traverse to prosecute those claims restricted to Group I. The provisional election is hereby affirmed. Accordingly, Applicants have canceled claims 14-18 and 24-29 for possible inclusion in a continuing application. Consideration of claims 1-13 and 19-23 is respectfully requested in view of the remarks set forth below.

Rejections under 35 U.S.C. § 102

The Examiner rejected claims 1, 3-7, 9, 10 and 19-21 under 35 U.S.C. § 102(e) as being anticipated by Hunsaker, (U.S. Pub. No. 2003-0037198). With specific regard to the independent claims, the Examiner stated:

As to claim 1, Hunsaker teaches a method of processing a request in a computer system, comprising the acts of: (a) initiating a read request from a requesting agent (read request from one of PCI/PCI-X devices 187), the requesting agent residing on a bus (bus 185), wherein the read request has an address corresponding to a memory location (in memory 140); (b) receiving the read request at a processor controller (part of bridge 190 “bridge 190 receives the read request from the PCI/PCI-X devices”); (c) sending the read request from the processor controller to an access

controller (Memory controller 130); (d) sending a deferred reply (split response) from the processor controller to the requesting agent when the processor controller is free to process the read request; (e) delivering data residing at the address corresponding to the memory location to the access controller (Memory controller 130); (f) delivering the data from the access controller to the processor controller (to bridge 190); and (g) delivering the data from the processor (from bridge 190) controller to the requesting agent.

As to claim 19, Hunsaker teaches a plurality of buses (buses 185, 195, 120, etc.); a memory system (System Memory 140) operably coupled to the plurality of buses (Fig. 1); and a processor controller (bridge 190, Memory controller 130 and I/O controller 150 in part or in combination) coupled to each of the plurality of buses and configured to simultaneously issue a deferred reply (split response) to a requesting device in response to receiving a read request from the requesting device and obtain the data corresponding to the read request from the memory system (Figs. 1-2 and paragraphs [0004], [0022] and [0023]).

Applicants respectfully traverse this rejection. Anticipation under Section 102 can be found only if a single reference shows exactly what is claimed. *Titanium Metals Corp. v. Banner*, 778 F.2d 775, 227 U.S.P.Q. 773 (Fed. Cir. 1985). For a prior art reference to anticipate under Section 102, every element of the claimed invention must be identically shown in a single reference. *In re Bond*, 910 F.2d 831, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990). To maintain a proper rejection under Section 102, a single reference must teach each and every element or step of the rejected claim. *Atlas Powder v. E.I. du Pont*, 750 F.2d 1569 (Fed. Cir. 1984). Thus, if the claims recite even one element not found in the cited reference, the reference does not anticipate the claimed invention.

Independent claim 1 recites “sending a deferred reply from the processor controller to the requesting agent when the processor controller is free to process the read request, *wherein the deferred reply is sent before data corresponding to the read request is delivered to the access controller.*” Emphasis added. Independent claim 19 recites a system comprising “a processor

controller coupled to each of the plurality of buses and configured to *simultaneously issue a deferred reply* to a requesting device in response to receiving a read request from the requesting device *and obtain the data* corresponding to the read request from the memory system.”

Emphasis added. As discussed further below, the Hunsaker reference does not disclose either of these features.

As discussed in the present application, in certain instances, a host controller is unable to immediately process a request received from a requesting agent. Page 10, line 13. In this case, rather than returning the requested data, the host controller defers the cycle, freeing it from the bus and indicating that the cycle will be completed at a later time. Page 10, lines 15-17. As discussed in the present application, in accordance with conventional techniques, if a request is received at a host controller and the request is unable to be processed immediately, the host controller defers the cycle to the requesting agent and retrieves the data from memory. Page 10, lines 15-17. Once the data is retrieved from the memory, it is temporarily restored in the host controller, and the host controller issues a “deferred reply” to the requesting bus. *See* page 10, lines 17-19 and page 11, lines 13-15.

As will be appreciated by those skilled in the art, a “deferred reply” is generally issued after the data is retrieved from memory and when the host controller and requesting bus are ready to complete the transaction. Once the bus is able to handle the data and deliver it to the requesting agent, the requested data is delivered to the requesting agent. Page 11, lines 14-15. Because of processor specific timing requirements and the associated architectural protocol, once a deferred reply is issued, a host controller waits some period of time after the issuance of the deferred reply until the appropriate data can be sent from the host controller to the requesting agent. Page 10, lines 19-23. Accordingly, by issuing the deferred reply only after data has been

retrieved from memory and delivered to the host controller, cycle time may be added to the processing of the request due to the processor specific timing requirements and the associated architectural protocol.

As will be appreciated by those skilled in the art, the Hunsaker reference simply describes the conventional techniques discussed above. Specifically, the Hunsaker reference discloses receiving a read request at a PCI bridge and issuing a delayed transaction or receiving a read request at a PCI-X bridge and issuing a split transaction. Paragraph [0004], lines 1-5. When a PCI bridge receives the read data from another device, it is stored in a transaction buffer in the PCI bridge. Paragraph [0004], lines 8-9. Once the same read request is resent by the requesting agent and received at the PCI bridge, the data is delivered to the requesting agent. Paragraph [0004], lines 8-10. Similarly, a PCI-X split transaction occurs when an I/O device issues an initial read request. Paragraph [0004], lines 9-11. Since the PCI-X bridge does not have the read data, it terminates the transaction with a split response indicating that the bridge has accepted the read request and will later provide the I/O device with the read completion data. Paragraph [0004], lines 11-15. When the bridge receives the read completion data, it stores it in a delayed transaction buffer. Paragraph [0004], lines 15-17. The bridge *then* sends the data to the I/O device as a split completion transaction. Paragraph [0004], lines 17-18.

In contrast to the conventional techniques described in the present application and disclosed in the Hunsaker reference, in accordance with an embodiment of the present invention, a more efficient method of processing requests by issuing the deferred reply to the requesting bus when the processor controller is free to process the request, regardless of whether the data has been retrieved from the memory and delivered to the host controller. Page 11, line 22 – page 12, line 1. Thus, once the requesting bus has enough bandwidth to handle the request, the host controller

issues the deferred reply. Page 12, lines 1-3. A deferred reply may be issued as early as the clock cycle immediately subsequent to the clock cycle in which the request was originally deferred. Page 12, lines 3-4. By initiating the deferred reply immediately upon the availability of the processor controller to process the request, the latent clock cycles associated with the issuance of the deferred reply and defined by the system protocol can be used in transferring the data from memory to the host controller. Page 12, lines 6-9. Accordingly, once the wait period has expired, the data may be waiting in the host controller for immediate delivery onto the requesting bus. Page 12, lines 9-12. By using the natural delays associated with standard protocols to carry out tasks (previously performed in series, as described above with reference to conventional techniques)) in parallel, there may be a reduction in the latency associated with processing the request. Page 12, lines 12-14; *See* Fig. 3.

The Examiner appears to be asserting that the split transaction of the PCI-X bus described in the Hunsaker reference indicates the same parallel processing described above. In other words, it appears that the Examiner believes that the conventional split transaction associated with PCI-X architectures provides for issuing a deferred reply from the host controller to the requesting agent, before the corresponding data is delivered from the memory to the host controller, as in claim 1, and simultaneously issuing the deferred reply to the requesting agent and obtaining data from the memory, as in claim 19. Applicants respectfully traverse this understanding.

As will be appreciated by those skilled in the art, the “split transaction” associated with the PCI-X architecture includes a “split response” and a “split completion.” *See e.g.*, Hunsaker, paragraph [0004], lines 11-19. Contrary to the Examiner’s assertions, at best, the split completion portion of the disclosed split transaction is analogous to the presently recited “deferred reply.” That is, the “split completion” transaction of Hunsaker is delivered from the PCI-X bridge to the

requesting agent to indicate that the PCI-X bridge is now ready to complete the transaction. As with the deferred reply, the “split completion” has a certain inherent delay associated therewith. As such, once the split completion is issued by the PCI-X bridge, the bridge waits a number of clock cycles (associated with processor specific timing requirements and the associated architectural protocol) before actually transferring the data from the PCI-X bridge to the requesting agent. As described in paragraph [0004] of the Hunsaker reference, the split completion transaction is issued only after the data is retrieved from the memory and temporarily stored in the PCI-X bridge. Accordingly, the conventional techniques described in the Hunsaker reference are subject to the same delays associated with the serial methods discussed above, wherein the deferred reply/ split completion transaction is issued only after data has been retrieved from memory and delivered to the host controller/ PCI-X bridge.

Accordingly, Applicants respectfully submit that the Hunsaker reference does not disclose the features recited in independent claims 1 and 19. Specifically, Applicants respectfully submit that the Hunsaker reference does not disclose “sending a deferred reply from the processor controller to the requesting agent when the processor controller is free to process the read request, *wherein the deferred reply is sent before data corresponding to the read request is delivered to the access controller,*” as recited in claim 1. Emphasis added. Further, Applicants respectfully submit that the Hunsaker reference does not disclose a system comprising “a processor controller coupled to each of the plurality of buses and configured to *simultaneously issue a deferred reply* to a requesting device in response to receiving a read request from the requesting device *and obtain the data* corresponding to the read request from the memory system,” as recited in claim 19. Emphasis added. Accordingly, the Hunsaker reference cannot possibly anticipate the recited subject matter. Thus, Applicants respectfully request withdrawal of the Examiner’s rejection and allowance of claims 1, 3-7, 9, 10 and 19-21.

Rejections under 35 U.S.C. § 103

The Examiner rejected claims 2 and 11-13 under 35 U.S.C. § 103(a) as being unpatentable over Hunsaker, (U.S. Pub. No. 2003-0037198), in view of Solomon (6,647,454).

With specific regard to independent claim 11, the Examiner stated:

As to claim 11, Hunsaker teaches a method of processing a request in a computer system, comprising the acts of : (a) sending a request from a processor controller (part of bridge 190 “bridge 190 receives the read request from the PCI/PCI-X devices”) to an access controller (Memory controller 130) the request originating from an agent; and (b) sending a deferred reply (split response) from the processor controller to the agent after step (a) (Figs. 1-2 and paragraphs [0004], [0023]) However, Hunsaker does not explicitly disclose sending the request is performed in a first clock cycle and sending the deferred reply is performed in a second clock cycle, wherein the second clock cycle being immediately subsequent to the first clock cycle. Solomon teaches sending the request is performed in a first clock cycle (phase 1) and sending the deferred reply is performed in the second clock cycle subsequent to the first clock cycle (Fig. 2 and col. 2, lines 47-61). It would have been obvious to one of ordinary skill in the art at the time the invention was made to implement sending the request is performed in a first clock cycle and sending the deferred reply is performed in a second clock cycle immediately subsequent to the first clock cycle as taught by Solomon in the system of Hunsaker to reduce wait stated in processing split requests (col. 1, lines 11-23).

Applicants respectfully traverse this rejection. The burden of establishing a *prima facie* case of obviousness falls on the Examiner. *Ex parte Wolters and Kuypers*, 214 U.S.P.Q. 735 (PTO Bd. App. 1979). Obviousness cannot be established by combining or modifying the teachings of the prior art to produce the claimed invention absent some teaching or suggestion supporting the combination or modification. *See ACS Hospital Systems, Inc. v. Montefiore Hospital*, 732 F.2d 1572, 1577, 221 U.S.P.Q. 929, 933 (Fed. Cir. 1984). Accordingly, to establish a *prima facie* case, the Examiner must not only show that the combination includes *all* of the claimed elements, but also a convincing line of reason as to why one of ordinary skill in

the art would have found the claimed invention to have been obvious in light of the teachings of the references. *Ex parte Clapp*, 227 U.S.P.Q. 972 (B.P.A.I. 1985).

Independent claim 11 recites “sending a request from a processor controller to an access controller on a first clock cycle,” and “sending a deferred reply from the processor controller to the agent on a second clock cycle, the second clock cycle being immediately subsequent to the first clock cycle.” As recognized by the Examiner, the Hunsaker reference does not disclose receiving a request at a host controller and immediately sending a deferred reply. As discussed above with regard to the rejections under 35 U.S.C. § 102, the Hunsaker reference discloses waiting for the data to be retrieved from the memory before issuing a split completion transaction (deferred reply). Contrary to the Examiner’s assertion, the Solomon reference does not cure this deficiency. As with the Hunsaker reference, the Solomon reference merely discusses inherent aspects of a split transaction, as related to a PCI bus. The Examiner appears to assert that the “split response termination” transaction is the same as the presently recited “deferred reply.” As will be appreciated by those skilled in the art, this correlation is not accurate. The techniques described in the Solomon reference will still be subject to the delays described above. That is, once the data is retrieved from memory, a deferred reply or split completion will be sent to the requesting agent. Thus, the Solomon reference does not disclose sending a deferred reply immediately after requesting the data from memory.

Because neither of the references alone or in combination discloses each of the features recited in independent claim 11, the cited combination cannot possibly render the recited subject matter obvious. Accordingly, Applicants respectfully request withdrawal of the Examiner’s rejection and allowance of claims 2 and 11-13.

The Examiner rejected claim 8, which is dependent on claim 1, under 35 U.S.C. § 103(a) as being unpatentable over Hunsaker, (U.S. Pub. No. 2003-0037198), in view of Ajanovic et al. (U.S. Pat. No. 5,761,444). Further, the Examiner rejected claim 22, which is dependent on claim 19, under 35 U.S.C. § 103(a) as being unpatentable over Hunsaker, (U.S. Pub. No. 2003-0037198), in view of Jayakumar et al. (U.S. Pat. No. 6,012,118). Finally, the Examiner rejected claim 23, which is dependent on claim 19, under 35 U.S.C. § 103(a) as being unpatentable over Hunsaker, (U.S. Pub. No. 2003-0037198), in view of Stallmo et al. (U.S. Pat. No. 5,613,059).

For the reasons set forth above with regard to the independent claims, Applicants respectfully submit that claims 8, 22 and 23 are also allowable. Applicants note that none of the cited references cure the deficiencies of the Hunsaker reference. Because none of the references alone or in combination disclose the recited features, the references cannot possibly render the recited subject matter obvious. Accordingly, Applicants respectfully request withdrawal of the Examiner's rejection and allowance of claims 8, 22 and 23.

New Claims

New claims 30-34 have been added. Applicants respectfully submit that no new subject matter has been added through the addition of the new claims. Independent claim 30 recites a system comprising "a processor controller coupled between the memory system and the requesting agent and configured to send a deferred reply from the processor controller to the requesting agent when the processor controller is free to process the read request, regardless of whether data corresponding to the read request has been delivered from the memory system to the processor controller." For at least the reasons set forth above, it is clear that none of the prior art references

disclose a system configured to send a deferred reply, regardless of whether the data has been delivered from the memory to the processor controller. Accordingly, Applicants respectfully submit that for at least this reason, claims 30-34 are allowable over the art of record.

Conclusion

In view of the remarks and amendments set forth above, Applicants respectfully request allowance of claims 1-13, 19-23 and 30-34. If the Examiner believes that a telephonic interview will help speed this application toward issuance, the Examiner is invited to contact the undersigned at the telephone listed below.

Respectfully submitted,

Date: June 30, 2004



Robert A. Manware
Reg. No. 48,758
(281) 970-4545

CORRESPONDENCE ADDRESS
HEWLETT-PACKARD COMPANY
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400